

Amendments to the Specification:

Please replace the paragraph beginning on page 1, line 7 with the following amended paragraph:

It has become popular to provide a guaranteed data communication level through a particular data network. Termed "Quality of Service" (QoS), it is a guarantee that end-to-end quality characteristics of data communication, such as latency, data loss, etc., will not exceed a specified level. There is an increasing demand on current networks to deliver QoS for such applications as interactive video communication and high quality video distribution. QoS requires a ~~sender of data~~ ~~network~~ to allocate enough network resource (e.g., bandwidth capacity) end to end for the application. Multi Protocol Label Switching (MPLS) is one technique that allows dedicated paths to be set up across a core network to achieve QoS. MPLS gives each packet entering a network a label or identifier that contains information indicative of the destination. However, as a network grows, the number of MPLS paths will increase significantly (on the order of n^2 , where n is the number of paths). Since MPLS paths, and the resource of each path, are preassigned, increasing the number of ~~paths~~ ~~nodes~~ will tend to decrease the resource that can be allocated to each path, limiting the capacity of each path. (As used herein, "resource" is the data handling or communicating capability of a path.) Also, if the paths are created each time a request for a data transfer through the core network is made, the job of path creation would become huge, and the computing power at each node of the core network would need to be immense.

Please replace the two paragraphs beginning on page 2, line 9 with the following amended paragraphs:

Broadly, according to the present invention, a number of relay nodes, such as ATM switches, routers, or similar devices, are interconnected by network links to form a core network for communicating data. The core network is coupled to and interconnects data transfer entities, such as sub-networks or data processing systems, by gateway elements.

Communicatively coupled to the gateway elements is a management system that operates to allocate to the ~~gateways~~ ~~gateway~~ portions of the data communicating resources of predetermined ones of the network links, thereby providing each of the gateway element routes through the core network from each gateway element to other of the gateway elements.

Data transfers across the core network, according to the present invention, from one data transfer entity coupled to one gateway element to another data transfer entity coupled to another gateway element begins

with a request to the associated gateway element with specifics as to the QoS resources needed for the data transfer. The gateway element that will be doing the sending (the “sending gateway”) will check to see if the resources allotted to it are sufficient to meet the needs of the request. If so, the requested data transfer is granted, and the sending gateway element performs the necessary packet adding for a QoS transfer as is conventional. If, however, the resources available to the sending gateway element are not sufficient, additional resources can be borrowed from other gateway element allocations by a request for “reconfiguration” to the management system. If sufficient resources can be re-allocated, the sending gateway element will permit the requested data transfer as before; if not, the request for a data transfer will be refused.

Please replace the paragraph beginning on page 3, line 30 with the following amended paragraph:

Figs. 8A and 8B illustrate **Gateway-Trunk Tables** maintained by each gateway element, containing information as to the data communicating capacity of each provisioned link available to such gateway element;

Please replace the paragraph beginning on page 5, line 3 with the following amended paragraph:

The access networks 12 will include various data processing elements and systems such as representatively shown by data processing systems 36 and 38. The data processing elements/systems within a particular access network may communicate among themselves via the particular access network, or the they may communicate with data processing elements/systems of other access networks through QoS paths established according to the present invention by requesting a data transfer of certain characteristics.

Please replace the two paragraphs beginning on page 5, line 14 with the following amended paragraphs:

Illustrated in Fig. 3 is the structure of the gateway 14 used to implement the individual gateways GA, GB, GC, and GD. As Fig. 3 shows, the gateway 14 (e.g., gateway GA) comprises a control unit 60 coupled to a gateway device 62 by a communicative coupling 64. The control section 60 provides the intelligence of the gateway 14, and includes a processor unit 66 and storage 68 – also coupled to one another by the coupling 64. The processor unit 66 may be implemented by a microprocessor element, or it may be a state machine or a multiple state

machine implementation. The storage 68 may be of any form of storage, including magnetic media, although preferably it is semiconductor memory.

The processor unit 66 manages information received from the NMS 30 in a number of tables, including gateway route table 70, a gateway trunk table 72, and a gateway trunk status table 74. In addition, the control unit 60 sets up and manages a gateway interface table 76 that identifies the different interfaces of the gateway device that connect to network links such as the access network associated with the gateway 14, and the link L (e.g., La or Lb) that couples the gateway 14 to the associated edge node 20 (Fig. 1).

Please replace the paragraph beginning on page 6, line 3 with the following amended paragraph:

During an initialization period, predetermined routes and pre-defined link resources (e.g., bandwidth, class, etc.) are created by the NMS 30. Fig. 4 is an illustrative example of possible one way routes through the core network 16 and assigned to the gateways GA and GD coupled to the edge nodes EA and ED, labeled R1, R2, ..., R5. For example, route R1 connects the edge nodes EA and EB to one another through the relay ~~node~~ nodes C1 and C2. Similarly, the gateway GA is assigned two separate routes R2 and R5 from edge node EA to edge node EC through (1) relay nodes C1 and C3 and C1 and C2. Routes R3 and R4 are assigned to the gateway GD. Similar one-way routes (not shown) would be assigned to the gateways GB and GC, but are not shown in order to not unduly complicate Fig. 4 and the following discussion.

Please replace the paragraph beginning on page 7, line 25 with the following amended paragraph (right parenthesis added):

From the viewpoint of the gateways to which the resources allotted the gateways, i.e., the provisioned links, are called the “Trunks” of that Gateway. Thus, referring to Fig. 4, in view of the fact that the distribution of resources by the TMS has resulted in providing the gateway GA with predefined routes R1, R2 and R5 containing core network links each provisioned with a particular bandwidth. Gateway GA is managing “Trunks” Aa, Ac, and Ae (which form route R1), Trunks Ad and Af (which, together with Trunk Aa form route R5) and Trunk Ag (which, together with Trunks Aa and Ac, form route R2). Gateway GD has similar Trunks. Figs. 8A and 8B are Gateway Trunk Tables 100A and 100D, identifying the Trunks managed by the Gateways GA and GD, respectively. The first (left-most) column 102 of each Table 100A, 100D

identifies the Trunks managed by that Gateway. The next column 104 provides the identification of the Provisioned Link that forms the Trunk, and the last column 106 specifies the bandwidth allocated the Trunk by the TMS 32.

Please replace the three paragraphs beginning on page 11, line 1 with the following amended paragraphs:

However, if the receiving gateway 12 finds that the interface does have the resources to handle the communication, in step 172, steps 172 and 174 are left in favor of step 176 where the receiving gateway sets up the necessary configuration to mark and forward the packets pursuant to the QoS path specified in the request. The receiving gateway 12 then will return an "approve" to the sending gateway. The sending gateway, in step 178, will also set up the necessary configuration for marking and forwarding packets according to the QoS requirements specified in the request by the requesting entity, and send an "approve" to the requesting entity. The procedure will then terminate with step 190. Thereafter, the requesting entity will begin the communication, send the packets necessary. When the requesting entity concludes its communication, it will send a "release" message to the associated gateway element 14. And in turn, ~~associate gateway element~~ the associated gateway element 14 will modify its Gateway Trunk Status Table 130 accordingly.

The major steps used for trunk reconfiguration process performed by the ~~receiving~~ sending gateway in step 180 of the path creation procedure 150 is illustrated in Fig. 13. Briefly, the trunk reconfiguration process, designated with the reference numeral 200 in Fig. 13, operates to locate extra resources that can be at least temporarily re-allocated to a trunk of a sending gateway looking for a QoS path in response to a request. Not having found a route with sufficient resources in any of the trunks managed by the sending gateway, an effort to find a route by re-allocating the available resources of one or more trunks is made by trunk reconfiguration.

Thus, as Fig. 13 shows, the trunk reconfiguration process 200 begins with step 202, after the sending gateway has determined in step 180 of the path creation procedure 150 that a reconfiguration attempt is warranted. Step 202 sees the sending gateway sending a request to TMS 32 (Fig. 1) that includes an identification of the trunk needing more resource (e.g., bandwidth), and the amount of bandwidth needed. In step 204, the TMS 32 will search the TMS Trunk Status Table 110 (Fig. 9) for alternate trunks (managed by other gateways 12) that share the common Provisioned Link with the trunk identified in the request. For example,

assume that the sending gateway 14 is gateway GA, and that after being unable to find a route with the resources necessary (i.e., bandwidth) for the QoS path requested by a sending entity of the access network associated with gateway GA (e.g., data processing system 36), the sending gateway initiates the trunk reconfiguration process 200. As Fig. 4, ~~and the gateway trunk table 80A, and the gateway trunk table 100A~~ show, all outbound routes from the gateway GA use link a (which, when allocated resource, becomes trunk Aa). This trunk, however, is not a candidate for reconfiguration because it is fully managed by the gateway GA; that is, no other gateway manages it. Fig. 4 shows that there are two routes to the edge node (EC) from the gateway element EA, R5 and R2. Continuing with Fig. 4, note that link c, which provides the trunk (when provisioned) Ac (Fig. 8A) managed by the gateway GA, is shared by trunk Dc managed by gateway GD. Network link d also shared by trunks Ad and Dd of routes R2 and R4. Thus, either of the trunks ~~Ad-Dc~~ and Dd are candidates for the trunk reconfiguration process, and one would be made known to TMS 32 by the request sent thereto from gateway GA.

Please replace the paragraph beginning on page 13, line 3 with the following amended paragraph:

If the gateway managing the alternate trunk determines, in step 212, that sufficient available resource does exist, step 212 will be left in favor of step 218, where the gateway managing the alternate trunk will send the TMS 32 an “approve,” and adjust its associated Gateway Trunk Status Table 130 (reducing the amount of resource indicated in both the “Bandwidth” and the “Available” columns for the alternate trunk). The TMS 32 will also modify its Gateway-TMS Trunk Status Table to reflect this re-allocation of resource.